

Back to [Before You Start](#); Next to [Advice for Advisors](#); Up to [How to Be a Good Graduate Student](#).

Doing Research

For many new graduate students, graduate school is unlike anything else they've done. Sometimes it's hard to know exactly what it is you're supposed to be learning. Yes, you have to complete a dissertation, but how do you start? What should you spend your time doing?

Graduate school is a very unstructured environment in most cases. Graduate students typically take nine hours or less of coursework per semester, especially after the second year. For many, the third year -- after coursework is largely finished and preliminary exams have been completed -- is a very difficult and stressful period. This is when you're supposed to find a thesis topic, if you're not one of the lucky few who has already found one. Once you do find a topic, you can expect two or more years until completion, with very few landmarks or milestones in sight.

The following sections talk about the day-to-day process of doing research, criticism and feedback, working on the thesis, and financial support for research.

The Daily Grind

Being a good researcher involves more than ``merely" coming up with brilliant ideas and implementing them. Most researchers spend the majority of their time reading papers, discussing ideas with colleagues, writing and revising papers, staring blankly into space -- and, of course, having brilliant ideas and implementing them.

A later section discusses the process and importance of becoming part of a larger research community, which is a critical aspect of being a successful researcher. This section contains ideas on keeping track of where you're going, and where you've been, with your research, staying motivated, and how to spend your time wisely.

Keeping a journal of your research activities and ideas is very useful. Write down speculations, interesting problems, possible solutions, random ideas, references to look up, notes on papers you've read, outlines of papers to write, and interesting quotes. Read back through it periodically. You'll notice that the bits of random thoughts start to come together and form a pattern, often turning into a research project or even a thesis topic. I was surprised, looking back through my journal as I was finishing up my thesis, how early and often similar ideas had cropped up in my thinking, and how they gradually evolved into a dissertation.

You'll have to read a lot of technical papers to become familiar with any field, and to stay current once you've caught up. You may find yourself spending over half of your time reading, especially at the beginning. This is normal. It's also normal to be overwhelmed by the amount of reading you think you ``should" do. Try to remember that it's impossible to read everything that might be relevant: instead, read selectively. When you first start reading up on a new field, ask your advisor or a fellow student what the most useful journals and conference proceedings are in your field, and ask for a list of seminal or ``classic" papers that you should definitely read. For AI researchers, a useful (if slightly outdated) starting point is Agre's (see [agre]) summary of basic AI references. Similar documents may exist for other research areas -- ask around. Start with these papers and the last few years of journals and proceedings.

Before bothering to read *any* paper, make sure it's worth it. Scan the title, then the abstract, then -- if you haven't completely lost interest already -- glance at the introduction and conclusions. (Of course, if your advisor tells you that this is an important paper, skip this preliminary step and jump right in!) Before you try to get all of the nitty-gritty details of the paper, skim the whole thing, and try to get a feel for the most important points. If it still seems worthwhile and relevant, go back and read the whole thing. Many people find it useful to take notes while they read. Even if you don't go back later and reread them, it helps to focus your attention and forces you to summarize as you read. And if

you do need to refresh your memory later, rereading your notes is much easier and faster than reading the whole paper.

A few other points to keep in mind as you read and evaluate papers:

1. Make sure the ideas described really worked (as opposed to just being theoretically valid, or tested on a few toy examples).
2. Try to get past buzzwords: they may sound good, but not mean much. Is there substance and an interesting idea underneath the jargon?
3. To really understand a paper, you have to understand the motivations for the problem posed, the choices made in finding a solution, the assumptions behind the solution, whether the assumptions are realistic and whether they can be removed without invalidating the approach, future directions for research, what was actually accomplished or implemented, the validity (or lack thereof) of the theoretical justifications or empirical demonstrations, and the potential for extending and scaling the algorithm up.

Keep the papers you read filed away so you can find them again later, and set up an online bibliography (BibTeX is a popular format, but anything consistent will do). I find it useful to add extra fields for keywords, the location of the paper (if you borrowed the reference from the library or a friend), and a short summary of particularly interesting papers. This bibliography will be useful for later reference, for writing your dissertation, and for sharing with other graduate students (and eventually, perhaps, advisees).

Staying Motivated

At times, particularly in the ``middle years," it can be very hard to maintain a positive attitude and stay motivated. Many graduate students suffer from insecurity, anxiety, and even boredom. First of all, realize that these are normal feelings. Try to find a sympathetic ear -- another graduate student, your advisor, or a friend outside of school. Next, try to identify why you're having trouble and identify concrete steps that you can take to improve the situation. To stay focused and motivated, it often helps to have organized activities to force you to manage your time and to do something every day. Setting up regular meetings with your advisor, attending seminars, or even extracurricular activities such as sports or music can help you to maintain a regular schedule.

Chapman (see [\[chapman\]](#)) enumerates a number of ``immobilizing shoulds" that can make you feel so guilty and unworthy that you stop making progress. Telling yourself that you **should** have a great topic, that you **should** finish in n years, that you **should** work 4, or 8, or 12 hours a day isn't helpful for most people. Be realistic about what you can accomplish, and try to concentrate on giving yourself positive feedback for tasks you do complete, instead of negative feedback for those you don't.

Setting daily, weekly, and monthly goals is a good idea, and works even better if you use a ``buddy system" where you and another student meet at regular intervals to review your progress. Try to find people to work with: doing research is much easier if you have someone to bounce ideas off of and to give you feedback.

Breaking down any project into smaller pieces is always a good tactic when things seem unmanageable. At the highest level, doing a master's project before diving into a Ph.D. dissertation is generally a good idea (and is mandatory at some schools). A master's gives you a chance to learn more about an area, do a smaller research project, and establish working relationships with your advisor and fellow students.

The divide-and-conquer strategy works on a day-to-day level as well. Instead of writing an entire thesis, focus on the goal of writing a chapter, section, or outline. Instead of implementing a large system, break off pieces and implement one module at a time. Identify tasks that you can do in an hour or less; then you can come up with a realistic daily

schedule. If you have doubts, don't let them stop you from accomplishing something -- take it one day at a time. Remember, every task you complete gets you closer to finishing.

Getting to the Thesis

The hardest part of getting a Ph.D. is, of course, writing the dissertation. The process of finding a thesis topic, doing the research, and writing the thesis is different from anything most students have done before. If you have a good advisor and support network, you'll be able to get advice and help in setting directions and goals. If not, you may need to be more independent. If this is the case, don't just isolate yourself from the world: try to go out and find the resources and support you need from professors, other graduate students, mailing lists, friends, family, and publications like this one.

Finding an Advisor

Finding the right advisor can help you immeasurably in successfully completing a thesis. You should ideally have selected the schools you applied to by identifying faculty members you'd like to work with. If not, start looking around as early as possible. Of course, the ideal advisor will be in the area you're interested in working in, and will actively be doing high-quality research and be involved in and respected by the research community.

Read research summaries by faculty members (which are usually published by the department), go to talks they give, and attend or audit courses given by professors you might be interested in working with. Talk to other graduate students and recent graduates. Ask them how their relationships with their advisors are/were, how quickly the advisor's students graduate, and how successful (well recognized, high-quality) their research is. What kinds of relationships do they have -- frequent interactions, collaborative work, encouraging independence? handing out topics or helping students to create individual research areas, or a more hands-off style?

Other things to find out about potential advisors:

1. What is the average time their Ph.D. students take to finish their degrees? What is the dropout rate for their students?
2. How long have they been on the faculty? There are advantages and disadvantages to being one of the first members of a new research group. On the positive side, you often have more freedom to choose your research topic and to influence the direction of the group's research. On the negative side, you may be more isolated (since there won't be older graduate students in the group), and your advisor won't have as much experience.

A good advisor will serve as a mentor as well as a source of technical assistance. A mentor should provide, or help you to find, the resources you need (financial, equipment, and psychological support); introduce you and promote your work to important people in your field; encourage your own interests, rather than promoting their own; be available to give you advice on the direction of your thesis and your career; and help you to find a job when you finish. They should help you to set and achieve long-term and short-term goals.

Once you identify one or more potential advisors, get to know them. Introduce yourself and describe the area you're interested in. Attend their research group meetings if they hold them regularly. Give them a copy of a research proposal if you have a good idea of what you want to work on, and ask for comments. Ask whether they have any TA or RA positions available, or if there are any ongoing research projects that you could get involved with. Read their published papers, and the work of their students. Drop by during office hours and ask questions or make comments. Offer to read drafts of papers -- and do more than just proofread (see the section on feedback).

The type of relationship that each student needs with an advisor will be different. Some students prefer to be given more direction, to have frequent contact, and to be "checked up on." Others are more independent. Some may need

contact but be self-conscious about asking for it. Other things that vary include what kinds of feedback is preferred (lots of ``random" ideas vs. very directed feedback (pointers)), working individually vs. in groups, working on an established research project vs. a new, independent effort; working in the same area as your advisor or doing an ``outside" thesis.

You may find that your thesis advisor doesn't always give you all of the mentoring that you need. Multiple mentors are common and useful; they may include other faculty members in your department or elsewhere, senior graduate students, or other colleagues (see the section on networking). You may want to seriously consider changing thesis advisors if your advisor is inaccessible or disinterested, gives you only negative feedback, doesn't have the technical background to advise you on your thesis, or harasses you (see the section on women).

The most important thing is to ask for (i.e., demand politely) what you need.

Finding a Thesis Topic

Doing a master's project is often a good idea (and is required by some schools). Although choosing an appropriately scaled-down topic may be difficult, having the ideal topic is also less important, since you will have the chance to move on after only a year or so. If you have a good idea of what you want to do your Ph.D. dissertation on, choosing a master's project that will lead into the dissertation is wise: you will get a head start on the Ph.D., or may decide that you're not interested in pursuing the topic after all (saving yourself a lot of work and grief farther down the road).

A good source of ideas for master's projects (and sometimes for dissertation topics) is the future work section of papers you're interested in. Try developing and implementing an extension to an existing system or technique.

Generally speaking, a good Ph.D. thesis topic is interesting to you, to your advisor, and to the research community. As with many aspects of graduate school, the balance you find will depend at least in part on the relationship you have with your advisor. Some professors have well defined long-term research programs and expect their students to contribute directly to this program. Others have much looser, but still related ongoing projects. Still others will take on anyone with an interesting idea, and may have a broad range of interesting ideas to offer their students. Be wary of the advisor who seems willing to let you pursue any research direction at all. You probably won't get the technical support you need, and they may lose interest in you when the next graduate student with a neat idea comes along.

If you pick a topic that you're not truly interested in simply because it's your advisor's pet area, it will be difficult to stay focused and motivated -- and you may be left hanging if your advisor moves on to a different research area before you finish. The same is true for choosing a topic because of its marketability: if you're not personally excited about the topic, you'll have a harder time finishing and a harder time convincing other people that your research is interesting. Besides, markets change more quickly than most people finish dissertations.

In order to do original research, you must be aware of ongoing research in your field. Most students spend up to a year reading and studying current research to identify important open problems. However, you'll never be able to read everything that might be relevant -- and new work is always being published.

Try to become aware and stay aware of directly related research -- but if you see new work that seems to be doing exactly what you're working on, don't panic. It's common for graduate students to see a related piece of work and think that their topic is ruined. If this happens to you, reread the paper several times to get a good understanding of what they've really been accomplished. Show the paper to your advisor or someone else who's familiar with your topic and whose opinions you respect. Introduce yourself to the author at a conference or by e-mail, and tell them about your work. By starting a dialogue, you will usually find that their work isn't quite the same, and that there are still directions open to you. You may even end up collaborating with them. Good researchers welcome the opportunity to interact and collaborate with someone who's interested in the same problems they are.

To finish quickly, it's usually best to pick a narrow, well defined topic. The downside of this approach is that it may not be as exciting to you or to the research community. If you're more of a risk-taker, choose a topic that branches out in a new direction. The danger here is that it can be difficult to carefully define the problem, and to evaluate the solution you develop. If you have a topic like this, it helps a lot to have an advisor or mentor who is good at helping you to focus and who can help you maintain a reasonably rigorous approach to the problem.

In the extreme case, if your topic is so out of the ordinary that it's unrelated to anything else, you may have difficulty convincing people it's worthwhile. Truly innovative research is, of course, exciting and often pays back in recognition from the research community -- or you could just be out in left field. If you have a far-out topic, be sure that people are actually **interested** in it, or you'll never be able to ``sell" it later, and will probably have trouble getting your work published and finding a job. In addition, it will be hard to find colleagues who are interested in the same problems and who can give you advice and feedback.

In any case, a good topic will address important issues. You should be trying to solve a real problem, not a toy problem (or worse yet, no problem at all); you should have solid theoretical work, good empirical results or, preferably, both; and the topic will be connected to -- but not be a simple variation on or extension of -- existing research. It will also be significant yet manageable. Finding the right size problem can be difficult. One good way of identifying the right size is to read other dissertations. It's also useful to have what Chapman [\[chapman\]](#) calls a ``telescoping organization" -- a central problem that's solvable and acceptable, with extensions and additions that are ``successively riskier and that will make the thesis more exciting." If the gee-whiz additions don't pan out, you'll still have a solid result.

Remember that a thesis is only a few years of your work, and that -- if all goes well -- your research career will continue for another 30 or 40. Don't be afraid to leave part of the problem for future work, and don't compare yourself to senior researchers who have years of work and publications to show for it. (On the other hand, if you identify too much future work, your thesis won't look very exciting by comparison.) Graduate students often pick overly ambitious topics (in theory, your advisor will help you to identify a realistic size problem). Don't overestimate what other people have done. Learn to read between the lines of grandiose claims (something else a good advisor will help you to do).

Some schools may require that you write a thesis proposal. Even if they don't, this is a good first step to take. It forces you to define the problem, outline possible solutions, and identify evaluation criteria; and it will help you to get useful feedback from your advisor and other colleagues. Writing a good thesis proposal will take up to several months, depending on how much background work and thinking you've already done in the process of choosing the topic.

The proposal should provide a foundation for the dissertation. First, you must circumscribe the problem and argue convincingly that it needs to be solved, and that you have a methodology for solving it. You must identify and discuss related work: has this problem been addressed before? What are the shortcomings of existing work in the area, and how will your approach differ from and be an improvement over these methods?

Present your ideas for solving the problem in as much detail as possible, and give a detailed plan of the remaining research to be done. The proposal should include, or be structured as, a rough outline of the thesis itself. In fact, unless your final topic differs significantly from your proposed topic (which many do), you may be able to reuse parts of the proposal in the thesis.

You will probably have to take an oral exam in which you present and/or answer questions about your proposal. Be sure that your committee members are as familiar as possible with your work beforehand. Give them copies of the proposal, and talk to them about it. During the exam, don't panic if you don't know the answer to a question. Simply say, ``I'm not sure" and then do your best to analyze the question and present possible answers. Your examining committee wants to see your analytical skills, not just hear canned answers to questions you were expecting. Give a practice talk to other students and faculty members. Remember: you know more about your thesis topic than your committee; you're teaching **them** something for a change.

Writing the Thesis

Graduate students often think that the thesis happens in two distinct phases: doing the research, and writing the dissertation. This may be the case for some students, but more often, these phases overlap and interact with one another. Sometimes it's difficult to formalize an idea well enough to test and prove it until you've written it up; the results of your tests often require you to make changes that mean that you have to go back and rewrite parts of the thesis; and the process of developing and testing your ideas is almost never complete (there's always more that you *could* do) so that many graduate students end up "doing research" right up until the day or two before the thesis is turned in.

The divide-and-conquer approach works as well for writing as it does for research. A problem that many graduate students face is that their only goal seems to be "finish the thesis." It is essential that you break this down into manageable stages, both in terms of doing the research and when writing the thesis. Tasks that you can finish in a week, a day, or even as little as half an hour are much more realistic goals. Try to come up with a range of tasks, both in terms of duration and difficulty. That way, on days when you feel energetic and enthusiastic, you can sink your teeth into a solid problem, but on days when you're run-down and unmotivated, you can at least accomplish a few small tasks and get them off your queue.

It also helps to start writing at a coarse granularity and successively refine your thesis. Don't sit down and try to start writing the entire thesis from beginning to end. First jot down notes on what you want to cover; then organize these into an outline (which will probably change as you progress in your research and writing). Start drafting sections, beginning with those you're most confident about. Don't feel obligated to write it perfectly the first time: if you can't get a paragraph or phrase right, just write *something* (a rough cut, a note to yourself, a list of bulleted points) and move on. You can always come back to the hard parts later; the important thing is to make steady progress.

When writing a thesis, or any technical paper, realize that your audience is almost guaranteed to be less familiar with your subject than you are. Explain your motivations, goals, and methodology clearly. Be repetitive without being boring, by presenting your ideas at several levels of abstraction, and by using examples to convey the ideas in a different way.

Having a "writing buddy" is a good idea. If they're working on their thesis at the same time, so much the better, but the most important thing is that they be willing to give you feedback on rough drafts, meet regularly to chart your progress and give you psychological support, and preferably that they be familiar enough with your field to understand and review your writing.

Getting Feedback

To be successful at research, it is essential that you learn to cope with criticism, and even that you actively seek it out. Learn to listen to valid, constructive criticism and to ignore destructive, pointless criticism (after finding any pearls of wisdom that may be buried in it).

In order to get feedback, you have to present your ideas. Write up what you're working on, even if you're not ready to write a full conference or journal paper, and show it to people. Even for pre-publishable papers, write carefully and clearly, to maximize your chances of getting useful comments (and of having people read what you wrote at all).

Give presentations at seminar series at your university, at conferences, and at other universities and research labs when you get the chance. Your advisor should help you find appropriate forums to present your work and ideas. Many fields have informal workshops that are ideal for presenting work in progress.

Attend conferences and talk about your research. When you meet someone new and they ask you what you're working

on, seize the opportunity. Don't just say "I'm doing my thesis on foobar applications of whatsis algorithms" -- tell them as much as they're willing to listen to. You should have 30-second, 2-minute, 5-minute and 10-minute summaries of your thesis ready at a moment's notice (but not memorized word-for-word; nobody wants to listen to a canned speech).

Talking to other people will help you to realize which aspects of your research are truly different and innovative, how your work fits into the current state of your field and where it's going, and which aspects of your work are harder to sell (and, therefore, which aspects you need to think more about justifying).

Giving feedback to other students and colleagues is useful for many reasons. First, it helps you to polish your critical skills, which are helpful both in understanding other people's work and in evaluating your own. Second, it helps you to build a network of people who will be your colleagues for years to come. Finally, if you give useful feedback, those people will be more likely to make an effort to do the same for you.

It will be helpful (to you and to the person whose paper you're reviewing) to organize comments on a paper in descending order of abstraction: high-level content-oriented comments, mid-level stylistic and presentation comments, and low-level nitpicky comments on syntax and grammar. Try to keep your comments constructive ("this would read better if you defined X before introducing Y") rather than destructive ("this is nonsense").

You'll want to read a paper at least twice -- once to get the basic ideas, then a second time to mark down comments. High-level comments describing your overall impression of the paper, making suggestions for organization, presentation and alternative approaches to try, potential extensions, and relevant references are generally the most useful and the hardest to give. Low-level comments are more appropriate for a paper that is being submitted for publication than for an unpublished paper such as a proposal or description of preliminary research.

Getting Financial Support

Most graduate students (at least in the natural sciences) have a source of financial support that pays their tuition and a small living stipend. Although nobody ever got rich being a graduate student, you probably won't starve either. Sources of funding include fellowships (from NSF, universities, foundations, government agencies, and industry), employer support, research assistantships (i.e., money from a faculty member's research grant) and teaching assistantships.

Start looking for money early. Many schools arrange support in the form of an RA or TA position in the first year, but after that, you're on your own. Deadlines for applications vary, and if you miss one, you'll probably have to wait another year. After you apply, it can take six months or so to review the applications and several more months to actually start receiving money.

Ask faculty members (especially your advisor, who should be helping you to find support or providing support out of his or her grant money), department administrators, and fellow graduate students about available funding. Go to your university's fellowship office or its equivalent, and look through the listings in "The Annual Register of Grant Support", "The Grant Register", "The Chronicle of Higher Education", and "Foundation Grants to Individuals". Look into NSF grants (there are several different programs). Take advantage of your status as a woman or minority if you are one (this may be the only time when it actually is an advantage). Most universities have fellowship programs that may be administered through individual departments or may be campus-wide.

If you haven't yet begun actively doing research, getting an RA position from a faculty member may be a good way to become involved in a research project. Working on an existing research project by maintaining or developing hardware or software, writing reports, and running experiments will give you a feel for what it's like to do research -- and you may even find a thesis topic. Ask around to see what's available, and go talk to professors whose work you find interesting.

For a research grant or fellowship, you will probably have to write a proposal, so the more you've thought about potential thesis topics, the better off you'll be. You may need to tailor your proposal to the interests and needs of the particular funding agency or program you're applying to, but stick to something you know about and are sincerely interested in.

Write for a general audience, since the people reviewing your application may not be in the same field. Emphasize your goals and why the project you propose to work on is important. Talk as much as you can about how you're going to solve the problem, and be sure that your proposed solution will satisfy the goals you've set forth. Follow the rules for format, page layout and length, or your application may not even be reviewed.

Next to [Advice for Advisors](#).